

ABSTRACT OF THE DISCLOSURE

The present invention, in various embodiments, provides techniques for managing latencies in accessing memory of computer systems. In one embodiment, upon accessing the memory system for a piece of data used by a first process, a latency manager

- 5 determines the access time to acquire the piece of data in the memory system. The latency manager then compares the determined access time to a threshold. If the determined access time is greater than the threshold, the latency manager triggers an interrupt for the operating system to switch threads or processes so that execution of the first process is postponed and execution of a second process starts. Various embodiments include the
- 10 latency manager is polled for the access time when the processor is stalled, the latency manager triggers a process switch when a particular memory subsystem is accessed, etc.

2011012382-1